

Altricial self-organising information-processing systems *

Aaron Sloman

(<http://www.cs.bham.ac.uk/~axs/>)

Jackie Chappell

(<http://www.biosciences.bham.ac.uk/staff/staff.htm?ID=90>)

1 Introduction

It is often thought that there is one key design principle or at best a small set of design principles, underlying the success of biological organisms. Candidates include neural nets, 'swarm intelligence', evolutionary computation, dynamical systems, particular types of architecture or use of a powerful uniform learning mechanism, e.g. reinforcement learning. All of those support types of self-organising, self-modifying behaviours. But we are nowhere near understanding the full variety of powerful information-processing principles 'discovered' by evolution. By attending closely to the *diversity* of biological phenomena we may gain key insights into (a) how evolution happens, (b) what sorts of mechanisms, forms of representation, types of learning and development and types of architectures have evolved, (c) how to explain ill-understood aspects of human and animal intelligence, and (d) new useful mechanisms for artificial systems.

2 The precocial–altricial spectrum

Consider the relative influence of nature and nurture during development. The vast majority of species (e.g. grazing mammals, chickens, fish, reptiles, insects, ...) are 'precocial': the young are born or hatched relatively well-developed and competent, with most behaviours genetically pre-programmed, whereas 'altricial' species start physically helpless and generally incompetent, requiring a period of support, including feeding, by parents. Paradoxically, some of these e.g. humans, primates, hunting mammals and nest-building birds, exhibit cognitive capabilities of far greater sophistication in adult life than precocial species.

What can in principle be achieved by genetic pre-programming is shown by 'cathedrals' produced by termites, and by cognitive systems that are sufficiently powerful within hours of birth to enable animals such as deer, to stand up, run with the herd, find a nipple and suck. In contrast, the fact that some animals requiring more complex and varied skills as adults tend to be start life helpless and incompetent suggests that evolution discovered limits to pre-programming and added something else, something of great power, apparently required for human intelligence.

The two labels 'precocial' and 'altricial' suggest a simple dichotomy between species whose behaviours are all innate and species whose behaviours result from learning and development. This is an oversimplification: there is a spectrum of cases. At every stage, in all animals, there are combinations of capabilities determined to varying degrees by the genome, embryonic development, maturation and kinds of learning.

Precocial behaviours are largely developmentally fixed, but often allow calibration and gradual re-shaping through maturation or processes like reinforcement learning. Conversely, species labelled 'altricial' because individuals start helpless and under-developed, nevertheless have some well-developed precocial skills at birth (e.g. those related to suckling in mammals and begging in birds), and some developmentally fixed capabilities manifested later, e.g. flying at first attempt, and migration skills in altricial birds.

Known mechanisms for learning and self-organisation explain neither the genetically determined sophistication at the precocial end, nor the richness and diversity of achievements of individuals of the same species at the altricial end. In particular, nobody knows how to design a robot with the precocial capabilities of a new-born deer, and no known learning mechanisms could transform a helpless infant-like robot placed in an any country into a lively talkative child.

3 Altricial, self-bootstrapping, architectures

Analysis of nature/nurture trade-offs, and variation in requirements for 'adult' information processing systems, reveals a need for previously unnoticed varieties of designs for artificial self-organising systems. Application domains where tasks and environments are fairly static and machines need to be functional quickly, require precocial skills (possibly including some adaptation and self-calibration), whereas others require altricial capabilities, e.g. where tasks and environments vary widely and change in complex ways over time, and where machines need to learn how to cope without being sent for re-programming. Architectures, mechanisms, forms of representation, and types of learning may differ sharply between the two extremes. And the end results of altricial learning by the same initial architecture may differ widely.

Many species require rich cognitive structures closely adapted to complex features of the environment. Sometimes those requirements change rapidly e.g. because individuals migrate to new terrain, or because climate patterns or geological catastrophes produce rapid environmental changes, or because other species, whether prey or predators, learn new behaviours, or new varieties arrive from elsewhere. If requirements change too fast for natural selection to keep up, and too fast for the forms of self-modification produced by mechanisms like reinforcement learning, a more powerful form of learning is needed, as evolution seems to have 'discovered'.

Where learnt capabilities involve collaboration with conspecifics, rapid cultural changes can cause additional pressures favouring mechanisms capable of rapidly acquiring complex non-innate knowledge, including novel ontologies – as shown by very young human children picking up concepts their parents never had to learn at that age. Such mechanisms,

*Many thanks to Birmingham colleagues on the CoSy project
<http://www.cs.bham.ac.uk/research/projects/cosy/>

in turn can speed up cultural change: a form of positive feedback. A special case is language learning: where phonology, syntax and vocabulary learnt by a child born in one country may be very different from what the parents learnt as children in other countries. What could produce so much structural variation in knowledge and behaviours within a species? Perhaps a new type of self-bootstrapping information-processing architecture evolved to enhance and complement both innate mechanisms and slow forms of individual learning.

4 Towards altricial architectures

We conjecture that altricial bootstrapping mechanisms, instead of being driven only by reward and punishment, also spontaneously discover discrete, re-usable and (recursively) recombinable chunks of information. Perceptual or action patterns found during spontaneous play and exploration, are selected for storage according to very general criteria e.g. symmetry, frequency, and production of complex effects through simple actions. Stored chunks can be used as components of larger chunks, which can be used in still larger chunks, using syntactic combination mechanisms forming conjunctions, sequences, loops, and conditional tests, enabling larger behavioural units to be formed, explored, and if found ‘interesting’ stored as new units.

Such mechanisms might discover ever more complex re-usable structures, in percepts and in actions, and store them for future use, both separately and in varied combinations, starting from (implicit) innate knowledge not about the specific features of the environment but about generic (meta-level) features that can be instantiated in different ways in different environments. There would have to be innate mechanisms for combining structures to form new, more complex, concepts, actions, strategies, percepts. More specifically, altricial learning may be based on genetically determined mechanisms with:

- implicit meta-level knowledge about kinds of information chunks that might be learnt, including
 - perceptual chunks (using concepts of space and time)
 - action chunks (using a concept of causation)
- implicit meta-level knowledge about kinds of associations that might be perceived, and knowledge about how to investigate which are causal and which spurious
- mechanisms for combining old chunks into more complex wholes (e.g. complex goals, or action sequences)
- mechanisms for discovering new complex wholes that occur in the environment, including
 - enduring objects that have persistent features, parts and patterns of behaviour,
 - processes extended in time in which objects endure even when not perceived
 - more and more complex actions produced and controlled by the individual
- mechanisms for creating and manipulating hypothetical structures which might describe unobserved portions of reality or possible future complex actions
- mechanisms for deriving consequences from complex information structures and for comparing and selecting between complex structures with different consequences.

Variants of such altricial mechanisms tailored to communication, might support development of languages with combinatorial syntax and semantics.

In humans, and perhaps some other species, altricial capabilities that were originally *outwardly* directed (e.g. perceiving and acting on external objects and processes) might, after suitable architectural extensions, also be *inwardly* directed, allowing individuals to develop more and more complex chunks of information not only about the environment,

but also about their own internal processes of perception, reasoning, learning, problem solving, motivation, choosing, etc.

Ontologies used for such internal ‘meta-management’ could also be used in mechanisms for perceiving, reasoning about and behaving towards others (e.g. conspecifics, prey and predators). Both the inward-directed and outward-directed cases require *meta-semantic* competence: the ability to represent and reason about entities which themselves process information. Animals and machines with such mechanisms can, for example, try to produce, change or prevent beliefs, plans or desires in others.

In humans one aspect of growth of the architecture seems to be acquisition of new sub-ontologies, new forms of representation, new collections of skills required for particular domains, e.g. learning a new language, learning to read music and play an instrument, learning programming, learning academic disciplines, learning athletic or dancing skills, learning mathematics, or quantum physics. Later growth enriches the architecture by growing new links between such domains – including using some as ‘metaphors’ for others.

Syntactic operations in such an altricial learner could themselves be either precocial (genetically determined) or altricial, i.e. made of more basic building blocks that are assembled into larger units by learnt during ‘playful’ thinking.

Individuals with such architectures are not limited to combinations of action units available at birth, but acquire more complex chunks indexed by their preconditions and effects. Searching for a combination that solves a complex problem may be very much faster than if the search either had to use more primitive units or had to use gradual modification of existing units.

5 The sources of meaning

The existence of mainly ‘precocial’ species shows that sophisticated visual and other apparatus can develop without individual learning. This means that the semantic content of the information structures is somehow determined by unlearnt structures and how they are applied, refuting theories that require all symbol-users to base their concepts on ‘symbol-grounding’ using processes of abstraction from experiences of instances. Instead, meaning can be largely determined by formal structures that limit possible (e.g. Tarskian) models, combined with ‘symbol attachment’ to reduce residual ambiguity. This helps to explain how altricial systems can develop theories about the unknown and unobservable – as humans do, e.g. in science. See also <http://www.cs.bham.ac.uk/research/cogaff/talks/#meanings>

6 Conclusion

These ideas seem to be close to Piaget’s theories about a child’s construction of reality. Which chunks an altricial individual learns will be influenced by physical actions possible for its body, the environment and its affordances, and the individual’s history. These factors could produce different kinds of understanding and representation of space, time, motion, causality and social relations in different species, or in similar individuals in different environments.

If all this is correct, after evolution discovered how to make physical bodies that grow themselves, it discovered how to make virtual machines that grow themselves. Whether computers as we know them can provide the infrastructure for such systems is a separate question.

A draft longer paper on this topic is online here

<http://www.cs.bham.ac.uk/research/cogaff/altricial-precocial.pdf>